Non-Mean-Field Behavior of Realistic Spin Glasses

C. M. Newman*
newman@cims.nyu.edu
Courant Institute of Mathematical Sciences
New York University
New York, NY 10012, USA

D. L. Stein†
dls@physics.arizona.edu
Dept. of Physics
University of Arizona
Tucson, AZ 85721, USA

Abstract: We provide rigorous proofs which show that the main features of the Parisi solution of the Sherrington-Kirkpatrick spin glass are not valid for more realistic spin glass models in any dimension and at any temperature.

The theoretical perspective provided by the Parisi solution [17] of the infinite-ranged Sherrington-Kirkpatrick (SK) model [21] has dominated the spin glass literature over the past decade and a half. This is partly because it represents the only example of a reasonably complete thermodynamic solution to an interesting and nontrivial spin glass model, and partly because of the novel, and in some respects, spectacular, nature of the symmetry breaking displayed in the low-temperature phase. Its main qualitative features — the presence of (countably) many pure states, the non-self-averaging of their overlap distribution function, and the ultrametric organization of their overlaps, among others — have greatly influenced thinking about disordered and complex systems in general [1, 19]. A common working hypothesis is that the Parisi solution provides a theory of general spin glass models [1, 14, 19]. In particular, many authors have directly applied its features to the study of both short-ranged models and laboratory spin glasses [2, 3, 8, 11]. Support for this "SK picture" — that the main qualitative features of Parisi's solution survive in non-infinite-ranged models — comes from both analytical [9] and numerical [5, 20] work.

In this Letter, however, we prove that this approach is fundamentally flawed. That is, short-ranged models such as the standard nearest-neighbor Edwards-Anderson (EA) model [6] cannot have at *any* temperature all the basic features of the Parisi solution. Furthermore, most of our arguments rely on little more than the homogeneity properties of the disorder, and thus are applicable to more realistic spin glass models such as models with long-ranged couplings or diluted RKKY interactions [22].

*Partially supported by the National Science Foundation under grants DMS-92-09053 and 95-00868.
†Partially supported by the U.S. Department of Energy under grant DE-FG03-93ER25155.

We do not attempt to resolve in this paper the closely related issue of whether short-ranged spin glass models have many pure thermodynamic states at sufficiently high dimension and low temperature, or only a single pair. The latter conjecture arises from a droplet model [7] based on a scaling *ansatz* [12, 4, 7]. Rather, we assert that *if* there are many pure states, their structure and that of their overlaps cannot be that of the SK picture.

The SK picture. — The Parisi solution, when applied to the EA model at fixed temperature T, suggests that there exist two related quantities which are non-self-averaging (i.e., depending on the realization $\mathcal{J}$ of the couplings): (i) a state $\rho_{\mathcal{J}}(\sigma)$, which is a Gibbs probability measure (at temperature T) on the microscopic spin configurations σ on all of Z^d, and (ii) a Parisi order parameter distribution $P_{\mathcal{J}}(q)$, which is a probability measure on the interval $[-1,1]$ of possible overlap values. These two are related as follows: if one chooses σ and σ' from the product distribution $\rho_{\mathcal{J}}(\sigma)\rho_{\mathcal{J}}(\sigma')$, then the overlap

$$Q = \lim_{L\to\infty} |\Lambda_L|^{-1} \sum_{x\in\Lambda_L} \sigma_x \sigma'_x \tag{1}$$

has $P_{\mathcal{J}}$ as its probability distribution. Here $|\Lambda_L|$ is the volume of a cube Λ_L of side length L centered at the origin in d dimensions.

A crucial component of the SK picture is that the decomposition of $\rho_{\mathcal{J}}$ into pure states is *countable* (i.e., a sum rather than an integral):

$$\rho_{\mathcal{J}}(\sigma) = \sum_\alpha W_{\mathcal{J}}^\alpha \rho_{\mathcal{J}}^\alpha(\sigma)\ . \tag{2}$$

If σ is drawn from $\rho_{\mathcal{J}}^\alpha$ and σ' from $\rho_{\mathcal{J}}^\beta$, then the expression in Eq. (1) equals its thermal mean,

$$q_{\mathcal{J}}^{\alpha\beta} = \lim_{L\to\infty} |\Lambda_L|^{-1} \sum_{x\in\Lambda_L} \langle\sigma_x\rangle_\alpha \langle\sigma_x\rangle_\beta \quad . \tag{3}$$

Thus $P_{\mathcal{J}}$ is given by

$$P_{\mathcal{J}}(q) = \sum_{\alpha,\beta} W_{\mathcal{J}}^\alpha W_{\mathcal{J}}^\beta \delta(q - q_{\mathcal{J}}^{\alpha\beta}) \quad . \tag{4}$$

In the SK picture, the $W_{\mathcal{J}}^\alpha$'s and $q_{\mathcal{J}}^{\alpha\beta}$'s are non-self-averaging quantities, except for $\alpha = \beta$ or its global flip (where $q_{\mathcal{J}}^{\alpha\beta} = \pm q_{EA}$). The average $P(q)$ of $P_{\mathcal{J}}(q)$ over the disorder distribution ν of the couplings is a mixture of two delta-function components at $\pm q_{EA}$ and a continuous part between them.

The countability of the decomposition of Eq. (2) is also employed to obtain the often-used result (see, for example, Refs. [1, 13, 2]) that the free energies of the lowest-lying states are independent random variables with an exponential distribution.

Both $\rho_{\mathcal{J}}$ and $P_{\mathcal{J}}$ are infinite-volume quantities and so must be obtained by some kind of thermodynamic limit. Naively, one might simply fix $\mathcal{J}$ and attempt to take a sequence of increasing volumes with, say, periodic boundary conditions. However, we argued in a previous paper [15] that the existence of multiple pure states is inconsistent with the existence of such a limit for *fixed* $\mathcal{J}$. Instead, there would be chaotic size dependence, so that infinite-volume limits can be achieved only through coupling-*dependent* boundary conditions. We will

see below that, nonetheless, $P_{\mathcal{J}}$ (and $\rho_{\mathcal{J}}$) can be obtained by natural limit procedures which are coupling-independent and which imply translation invariance for $P_{\mathcal{J}}$ (and translation covariance for $\rho_{\mathcal{J}}$). We ask whether this is consistent with the SK picture, which requires the following properties of $P_{\mathcal{J}}$ and its average P:

1) $P_{\mathcal{J}}(q)$ is non-self-averaging.
2) $P_{\mathcal{J}}(q)$ is a sum of (infinitely many) delta-functions.
3) $P(q)$ has a continuous component (for all q between the delta-functions at $\pm q_{EA}$).

The answer is no; we will see that *translation invariance rules out non-self-averaging*. This in turn makes the absence of a continuous component in $P_{\mathcal{J}}$ inconsistent with its presence in P. We conclude that *property 1) is absent, and at most one of the remaining two properties of the SK picture can be valid for realistic spin glass models*. We will consider below the implications of this result for other important features of the SK picture, such as ultrametricity.

Construction of $\rho_{\mathcal{J}}$ and $P_{\mathcal{J}}$. — We first describe a limit procedure to obtain $P_{\mathcal{J}}$ which does not involve the prior construction of $\rho_{\mathcal{J}}$. Begin with the finite volume Gibbs distribution $\rho^{(L)}_{\mathcal{J}^{(L)}}$ on the spin configuration $\sigma^{(L)}$ in the cube, Λ_L, with periodic boundary conditions. Here $\mathcal{J}^{(L)}$ denotes the couplings restricted to Λ_L. Let $Q^{(L)}$ denote the overlap of $\sigma^{(L)}$ and a duplicate $\sigma'^{(L)}$:

$$Q^{(L)} = |\Lambda_L|^{-1} \sum_{x\in\Lambda^{(L)}} \sigma^{(L)}_x \sigma'^{(L)}_x \quad . \tag{5}$$

The distribution $P^{(L)}_{\mathcal{J}^{(L)}}$ for $Q^{(L)}$ is the finite volume Parisi overlap distribution function, whose average was studied numerically in Refs. [5, 20]. It was proved in Ref. [15] that in the SK model, non-self-averaging requires $P^{(L)}_{\mathcal{J}^{(L)}}$ to have chaotic L-dependence as $L \to \infty$ for fixed $\mathcal{J}$; a similar result was suggested, though not proved, for short-ranged spin glasses with many pure states. Because of this, we do not take a limit of $P^{(L)}_{\mathcal{J}^{(L)}}$ directly but rather of the *joint distribution* $\tilde{\mu}_L$ of $\mathcal{J}^{(L)}$ and $Q^{(L)}$. That is, by a compactness argument (which may require the use of a subsequence of L's) one has a limiting $\tilde{\mu}$, which is a probability measure on joint configurations $(\mathcal{J}, q)$ (q being a realization of Q) such that for any (nice) function f of *finitely* many couplings and of q, the average $\langle f\rangle$ for $\tilde{\mu}$ is the limit of the averages for $\tilde{\mu}_L$.

This gives us existence of a $\tilde{\mu}$, which is a joint distribution on the infinite-volume realizations of $\mathcal{J}$ and q. Its marginal distribution for $\mathcal{J}$ is the original disorder distribution ν, while its conditional distribution for q given $\mathcal{J}$ is what we denote $P_{\mathcal{J}}$. Because of the periodic boundary conditions, the marginal distribution (under $\tilde{\mu}_L$) of $J_1, \ldots, J_m, q$ is the same (for large L) as of $J^a_1, \ldots, J^a_m, q$ (where a is any lattice translation and $\mathcal{J}^a$ is the translated $\mathcal{J}$) and thus one has translation invariance of the limit measure $\tilde{\mu}$. Translation invariance here means that for any a, the shifted variables $\mathcal{J}^a$ together with Q have the same joint distribution as do the original $\mathcal{J}$ together with Q; because ν is in any case translation-invariant, this implies that $P_{\mathcal{J}} = P_{\mathcal{J}^a}$. In other words, the overlaps don't care about the choice of origin.

Our second procedure for obtaining $P_{\mathcal{J}}$ is first to construct $\rho_{\mathcal{J}}$ and then obtain $P_{\mathcal{J}}$ as the distribution of the Q given by Eq. (1). The construction of $\rho_{\mathcal{J}}$ is as follows. Let μ_L

be the joint distribution for $\mathcal{J}^{(L)}$ and $\sigma^{(L)}$ on the periodic cube Λ_L. Then by compactness arguments, some subsequence μ_L converges to a limiting joint distribution $\mu(\mathcal{J},\sigma)$. The resulting conditional distribution of σ given $\mathcal{J}$ is what we denote $\rho_{\mathcal{J}}(\sigma)$. μ will be translation-invariant (and $\rho_{\mathcal{J}}$ will be translation-covariant) because of the translation invariance (on the torus) of μ_L. Translation invariance means that the distribution μ for $(\mathcal{J},\sigma)$ is the same as for $(\mathcal{J}^a,\sigma^a)$ for any lattice vector a. In terms of $\rho_{\mathcal{J}}$, this means that $\rho_{\mathcal{J}^a}(\sigma)=\rho_{\mathcal{J}}(\sigma^{-a})$, so that, e.g., $\langle\sigma_x\rangle_{\mathcal{J}^a}=\langle\sigma_{x-a}\rangle_{\mathcal{J}}$; thus we say that $\rho_{\mathcal{J}}$ is translation-covariant rather than invariant.

Before pursuing the implications of translation invariance, we raise several questions related to our constructions. Could different subsequences of cubes yield different limits? We believe the answer is no, although we have no complete proof, because our procedure of considering *joint* distributions (for $\mathcal{J}$ and q or for $\mathcal{J}$ and σ) should eliminate the kind of chaotic volume dependence discussed in Ref. [15]. Could different deterministic boundary conditions yield different limits? Boundary conditions related to each other by partial or complete spin flips (e.g., periodic and antiperiodic) must have the same limiting joint distributions (by arguments similar to those used in Ref. [15]), but, in principle, unrelated boundary conditions such as periodic and free could have different limits. In practice, however, we expect that different sequences of deterministic boundary conditions would yield the same (translation-invariant) limit. Could the $P_{\mathcal{J}}$'s arising from our two constructions (one using $\rho_{\mathcal{J}}$ and one not) be different? We shall take as a working hypothesis that the two are the same, but see no compelling reason why that should be the case. Either way, since both $P_{\mathcal{J}}$'s are translation-invariant, neither one can be non-self-averaging, as we now explain.

Self-averaging of $P_{\mathcal{J}}(q)$. — To justify our claim that translation invariance of $P_{\mathcal{J}}(q)$ implies that it is self-averaging, take a (nice) function $f(q)$ (like q^k) and consider the function of $\mathcal{J}$, $\hat{f}(\mathcal{J})\equiv\int f(q)P_{\mathcal{J}}(q)\ dq$. By translation invariance, $\hat{f}(\mathcal{J})=\hat{f}(\mathcal{J}^a)$, but by the *translation-ergodicity* [23] of ν, any translation-invariant (measurable) function $\hat{f}(\mathcal{J})$ is equal to its $\mathcal{J}$-average, $\int\hat{f}(\mathcal{J})\nu(\mathcal{J})\ d\mathcal{J}$. Since this is true for all f's, it follows that $P_{\mathcal{J}}$ itself equals its $\mathcal{J}$-average.

We remark that the above discussion makes it clear that our claim is valid for any model involving disorder whose underlying distribution is (like ν) translation-invariant and translation-ergodic. For example, any analogue of the Parisi order parameter distribution for spin glass models with site-diluted RKKY interactions will also be self-averaging (if it is translation-invariant).

Because $P_{\mathcal{J}}$ is self-averaging, we are forced to the dichotomy that, for any temperature in any dimension, either $P(=P_{\mathcal{J}})$ is a sum of one or more δ-functions or else P has a continuous component. When there is a unique infinite-volume Gibbs state (e.g., in the paramagnetic phase) then of course $\rho_{\mathcal{J}}$ is that state and P is a single δ-function at $q=0$. If there were only two pure states (related by a global flip) [10], then P would simply be a sum of two δ-functions at $\pm q_{EA}$. But what if infinitely many pure states $\rho^\alpha_{\mathcal{J}}$ coexist in $\rho_{\mathcal{J}}$, with infinitely many overlap values $q^{\alpha\beta}_{\mathcal{J}}$? If the set of overlap values were *countably* infinite, then $P_{\mathcal{J}}$ would necessarily be a sum of δ-functions, but *the infinitely many locations (as well as the weights)*

would not depend on $\mathcal{J}$; we regard this possibility as implausible.

Thus we suggest that the most likely scenario for many coexisting pure states and overlaps is one where the countable decomposition Eq. (2) is replaced by an integral and where P has no δ-function components.

Ultrametricity. — We briefly turn to the question of whether ultrametricity of pure state overlaps [24] can survive in short-ranged spin glasses, given that $P_{\mathcal{J}}$ is self-averaged. Clearly, this type of nontrivial ultrametricity requires the existence of multiple pure states. As discussed above, we consider the case where $P(q)$ is continuous. We now demonstrate that such an overlap distribution cannot have an ultrametric structure, in the Parisi sense.

Let $\alpha, \beta, \gamma_1, \gamma_2 \ldots$ denote pure states randomly selected from the continuum of such states, and let their overlaps as usual be denoted $q^{\alpha\beta}$, etc. In the Parisi solution, these overlaps are such that, for any k, the two smallest of $q^{\alpha\beta}, q^{\alpha\gamma_k}$, and $q^{\beta\gamma_k}$ are equal. For nontrivial ultrametricity such as occurs in the Parisi solution, there would be positive probability that for some i and j the following two strict inequalities occur simultaneously:

$$q^{\alpha\gamma_i} < q^{\beta\gamma_i} \quad \text{and} \quad q^{\alpha\gamma_j} > q^{\beta\gamma_j} \quad . \tag{6}$$

If ultrametricity holds, then the first inequality requires that $q^{\alpha\gamma_i} = q^{\alpha\beta}$, while the second inequality requires that $q^{\beta\gamma_j} = q^{\alpha\beta}$. Thus $q^{\alpha\gamma_i} = q^{\beta\gamma_j}$. But because α, β, γ_i, and γ_j are chosen randomly and independently, the two variables $q^{\alpha\gamma_i} = q^{\beta\gamma_j}$ are also independent. Because each of these is chosen from a *continuous* distribution P, the probability that the two overlap values can be identical is zero, and we arrive at a contradiction.

The only way to avoid the contradiction is if the two strict inequalities in Eq. (6) *cannot* occur simultaneously. This means that either $q^{\alpha\gamma_k} \leq q^{\beta\gamma_k}$ for *every* k or vice-versa, which implies that the pure states can be ordered into a one-dimensional continuum, and the ultrametric structure resembles a comb rather than a more usual tree, such as appears in the SK picture.

As discussed in the previous section (see also the next section), self-averaging makes it implausible that the set of overlaps is countable. A countable set of overlaps would invalidate the above argument and possibly rescue ultrametricity, but at the cost of destroying anything resembling the Parisi solution.

Decomposition into pure states. — What is the nature of the decomposition of $\rho_{\mathcal{J}}$ into pure states? The SK picture prediction of a countably infinite sum as in Eq. (2) (with infinitely many $q_{\mathcal{J}}^{\alpha\beta}$'s) has been largely ruled out since the set of $q_{\mathcal{J}}^{\alpha\beta}$'s would be self-averaging, which seems unreasonable. Even if one were unwilling to rule out countability on that ground, there are other arguments, not presented here, which make that possibility even more unlikely. These arguments suggest that all $\rho_{\mathcal{J}}^{\alpha}$'s appearing in a countable decomposition would have the same even spin correlations. This certainly seems inconsistent with the expected presence of domain walls between pure states unrelated by a global spin flip. We conclude that in any reasonable scenario for $\rho_{\mathcal{J}}$, there should be at most one pair of pure states (related by a global spin flip) with *strictly* positive weight.

In other words, either (a) $\rho_{\mathcal{J}}$ is pure, or (b) it is a sum of two pure states related by a global flip, or (c) it is an integral over pure states with none having strictly positive weight,

or (d) it has one “special” pair of pure states with strictly positive weight and all the rest with zero weight. Case (a) occurs if the system is in a paramagnetic phase, or any other in which the EA order parameter is zero. Case (b) would occur according to the Fisher-Huse droplet picture [7], but could also occur if there existed multiple pure states not appearing in $\rho_{\mathcal{J}}$ (“weak Fisher-Huse”) [25]. Case (c) occurs if there are *uncountably* many pure states in the decomposition of $\rho_{\mathcal{J}}$, all with zero weight (“democratic multiplicity”). Case (d) (which we regard as unlikely) occurs when one pair of pure states partially dominates all others, but accounts for only part of the total weight (“dictatorial multiplicity”).

What is the relation between the nature of $P(= P_{\mathcal{J}})$ and the three (nontrivial) cases (b) – (d) discussed above? Clearly case (b) implies that P is a sum of two δ-functions at $\pm q_{EA}$, and no continuous part. If we assume in cases (c) and (d) that varying α and β through the continuous portion of the pure states yields a continuously varying $q_{\alpha\beta}$ (but see the next paragraph for a case where this assumption is violated), then it follows that case (c) corresponds to a P with *no* δ-functions while case (d) corresponds to a P with δ-functions at $\pm q_{EA}$ *and* a continuous part. This latter case is the P predicted by the Parisi solution, but note two crucial distinctions between case (d) and the SK picture: (i) there is self-averaging, so one already obtains the continuous part of P from a single realization $\mathcal{J}$, and (ii) the δ-functions at $\pm q_{EA}$ come from a *single* special pair of pure states — not from countably many $q^{\alpha\alpha}$'s. *Thus any numerical evidence in favor of such a distribution for* $[P_{\mathcal{J}}]_{\rm av}$ (as in Refs. [5, 20]) *is not evidence in favor of the SK picture* but rather, at most, supports dictatorial multiplicity.

We remark that a case of democratic multiplicity occurs in a solution for the ground state structure in a short-ranged, highly disordered spin glass model [16]. We argued there that below eight dimensions, there exists a single pair of ground states (case (b) above), while above eight, there are uncountably many. It is not hard to see that the $\rho_{\mathcal{J}}$ for $d > 8$ corresponds to case (c) above – the states are chosen by the flips of fair coins for all the trees in the invasion forest, so all have equal (zero) weight. It appears that here $P(q)$ is a δ-function at zero! So even for short-ranged spin glasses with $T > 0$, we cannot rule out the possibility of democratic multiplicity with such a P.

Discussion and Conclusions. — We have shown that in realistic spin glass models, and probably all non-infinite-ranged spin glasses, a natural construction leads to a Parisi overlap distribution $P_{\mathcal{J}}$ which is translation-invariant and hence self-averaging, unlike the Parisi solution of the SK model. Although our construction uses periodic boundary conditions, we believe that, in short-ranged (and probably all non-infinite-ranged) models, any choice of *coupling-independent* boundary conditions (e.g., free) should yield the same translation-invariant $P_{\mathcal{J}}$. Any restoration of the SK picture would require a non-translation-invariant $P_{\mathcal{J}}$ and we see no natural mechanism for obtaining one. Even were such a mechanism available, we note that the construction of our self-averaged $P_{\mathcal{J}}$ shows that the SK picture can be, at best, incomplete. Any theory of the thermodynamics of realistic spin glasses will likely differ considerably from the SK picture.

If many phases do exist in some dimension and below some temperature, we believe that the most reasonable possibility is (c) above, i.e., democratic multiplicity. If numerical

experiments find an order parameter distribution that looks similar to that of the Parisi solution, our arguments show that it should be interpreted within the context of possibility (d) above, i.e., dictatorial multiplicity, rather than as confirmation of the SK picture.

To summarize our results, we have ruled out non-self-averaging in an extremely large class of disordered systems, which include short-ranged and probably all non-infinite-ranged spin glasses. Non-self-averaging, and the other main consequences of the Parisi solution, including ultrametricity of pure state overlaps, appear to be confined to mean field models.

References

[1] K. Binder and A.P. Young, *Rev. Mod. Phys.* **58**, 801 (1986).

[2] J.P. Bouchaud, *J. Phys. I (France)* **2**, 1705 (1992).

[3] J.P. Bouchaud, E. Vincent, and J. Hammann, *J. Phys. I (France)* **4**, 139 (1994).

[4] A.J. Bray and M.A. Moore, in *Heidelberg Colloquium in Glassy Dynamics*, ed. J.L. van Hemmen and I. Morgenstern (Berlin: Springer-Verlag, 1987); A.J. Bray and M.A. Moore, *Phys. Rev. Lett.* **58**, 57 (1987).

[5] S. Caracciolo, G. Parisi, S. Patarnello, and N. Sourlas, *J. Phys. France* **51**, 1877 (1990).

[6] S. Edwards and P.W. Anderson, *J. Phys. F* **5**, 965 (1975).

[7] D.S. Fisher and D.A. Huse, *Phys. Rev. Lett.* **56**, 1601 (1986); D.S. Fisher and D.A. Huse, *Phys. Rev. B* **38**, 386 (1988).

[8] S. Franz and M. Mézard, *Physica A* **210**, 48 (1994).

[9] A. Georges, M. Mézard, and J.S. Yedidia, *Phys. Rev. Lett.* **64**, 2937 (1990).

[10] D.A. Huse and D.S. Fisher, *J. Phys. A* **20**, L997 (1987); D.S. Fisher and D.A. Huse, *J. Phys. A* **20**, L1005 (1987).

[11] M. Lederman, R. Orbach, J.M. Hamann, M. Ocio, and E. Vincent, *Phys. Rev. B* **44**, 7403 (1991).

[12] W.L. McMillan, *J. Phys. C* **17**, 3179 (1984).

[13] M. Mézard, G. Parisi, N. Sourlas, G. Toulouse, and M. Virasoro, *Phys. Rev. Lett.* **52**, 1156 (1984); *J. Phys. (Paris)* **45**, 843 (1984); A.P. Young, A.J. Bray, and A.M. Moore, *J. Phys. C* **17**, L149 (1984); **17**, L155 (1984); M. Mézard, G. Parisi, and M. Virasoro, *J. Phys. (Paris) Lett.* **46**, L217 (1985); B. Derrida and G. Toulouse, *J. Phys. (Paris) Lett.* **46**, L223 (1985).

[14] M. Mézard, G. Parisi, and M.A. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987.)

[15] C.M. Newman and D.L. Stein, *Phys. Rev. B* **46**, 973 (1992).

[16] C.M. Newman and D.L. Stein, *Phys. Rev. Lett.* **72**, 2286 (1994).

[17] G. Parisi, *Phys. Rev. Lett.* **43**, 1754 (1979).

[18] G. Parisi, *Phys. Rev. Lett.* **50**, 1946 (1983); A. Houghton, S. Jain, and A.P. Young, *J. Phys. C* **16**, L375 (1983).

[19] G. Parisi, *Physica A* **194**, 28 (1993).

[20] J.D. Reger, R.N. Bhatt, and A.P. Young, *Phys. Rev. Lett.* **64**, 1859 (1990).

[21] D. Sherrington and S. Kirkpatrick, *Phys. Rev. Lett.* **35**, 1972 (1975).

[22] For such models, however, there is no general theory of infinite volume states or their decompositions into pure phases.

[23] By translation-ergodicity, we simply mean that the spatial average of a random variable equals its expectation. This is equivalent to all translation-invariant random variables being constants.

[24] We caution that this should not be confused with ultrametric hopping dynamics between *metastable* states. For a discussion of how the latter can arise, see D.L. Stein and C.M. Newman, *Phys. Rev. E* **51**, 5228 (1995).

[25] This would be analogous to "weak uniqueness" for spin glasses, as discussed in M. Campanino, E. Olivieri, and A.C.D. van Enter, *Comm. Math. Phys.* **108**, 241 (1987).